

CLAIMS

What is claimed is:

1. A computer system, comprising:
 - a host processor, and
 - a network interface coupled to the host processor and to a network, the network interface comprising:
 - a first port that receives data from the host processor;
 - a second port that transmits data to the network;
 - a memory that stores data packets received by the first port, the memory being coupled to the first port and to the second port;
 - a control circuit that manages the memory as a plurality of queues having respective priorities, including logic to place a packet received from the host into one of the plurality of queues according to a quality of service parameter associated with the packet, and logic to service packets in the plurality of queues according to the respective priorities; and
 - logic to dynamically allocate space in said memory to the queues in the plurality of queues.
2. The computer system of claim 1, wherein the plurality of queues includes a higher priority queue, and a lower priority queue, and including a timeout timer coupled with the lower priority queue which is enabled if a packet is stored in the lower priority queue and expires after a timeout interval, and including logic to preempt the higher priority queue in favor of the lower priority queue if the timeout timer expires.
3. The computer system of claim 1, wherein the plurality of queues includes a higher priority queue, an intermediate priority queue, and a lower priority queue, and including
 - a first timeout timer coupled with the intermediate priority queue which is enabled if a packet is stored in the intermediate priority queue and expires after a first timeout interval, and
 - including logic to preempt the higher priority queue in favor of the intermediate priority queue if the first timeout timer expires; and

7 a second timeout timer coupled with the lower priority queue which is enabled if a packet
8 is stored in the lower priority queue and expires after a second timeout interval, and including
9 logic to preempt the higher priority queue and the intermediate priority queue in favor of the
10 lower priority queue if the second timeout timer expires.

1 4. The computer system of claim 1, wherein the plurality of queues includes a higher
2 priority queue, an intermediate priority queue, and a lower priority queue, and including
3 an first timeout timer coupled with the intermediate priority queue which is enabled if a
4 packet is stored in the intermediate priority queue and expires after a first timeout interval, and
5 including logic to preempt the higher priority queue in favor of the intermediate priority queue if
6 the first timeout timer expires;

7 a second timeout timer coupled with the lower priority queue which is enabled if a packet
8 is stored in the lower priority queue and expires after a second timeout interval, and including
9 logic to preempt the higher priority queue and the intermediate priority queue in favor of the
10 lower priority queue if the second timeout timer expires; and

11 logic to service the intermediate priority queue in favor of the lower priority queue if both
12 the first and second timeout timers expire.

1 5. The computer system of claim 1, further comprising logic in the network interface to
2 execute a security process on packets in one of the plurality of queues.

1 6. The computer system of claim 1, wherein the second port further comprises circuitry for
2 formatting packets according to a protocol compliant with an Ethernet protocol standard.

1 7. The computer system of claim 1, wherein the second port further comprises circuitry for
2 formatting packets according to a protocol compliant with an InfiniBand protocol standard.

1 8. The computer system of claim 1, wherein the packets include frame start headers, and
2 said quality of service parameters comprises codes in the frame start headers.

1 9. The computer system of claim 1, wherein said logic to dynamically allocate space in said
2 memory maintains a list of free buffers and a list of used buffers for each of the plurality of
3 queues.

1 10. The computer system of claim 1, wherein said memory comprises a single storage array.

1 11. The computer system of claim 1, wherein said logic to dynamically allocate space in said
2 memory, places packets downloaded to the memory in a plurality of buffers in non-contiguous
3 memory locations.

1 12. The computer system of claim 1, wherein said logic to dynamically allocate space in said
2 memory maintains a list of free buffers and a list of used buffers for each of the plurality of
3 queues, so that each virtual path has a free buffer list having a number of free buffers, and
4 includes logic which releases a used buffer to the free buffer list for a queue in the plurality of
5 queues having a smallest number of free buffers.

1 13. The computer system of claim 1, wherein said logic to dynamically allocate space in said
2 memory maintains a list of free buffers and a list of used buffers for each of the plurality of
3 queues, so that each virtual path has a free buffer list having a number of free buffers, and
4 includes logic which releases a used buffer to the free buffer list for a queue in the plurality of
5 queues having a largest amount of traffic.

1 14. The computer system of claim 1, wherein logic to dynamically allocate space in said
2 memory maintains a list of buffer descriptors for corresponding buffers in said memory, the
3 buffer descriptors including a parameter specifying a size of the corresponding buffer, and a
4 location of the corresponding buffer.

1 15. The computer system of claim 1, wherein logic to dynamically allocate space in said
2 memory maintains a list of buffer descriptors for corresponding buffers in said memory, the
3 buffer descriptors including a programmable parameter specifying a size of the corresponding
4 buffer.

1 16. The computer system of claim 1, wherein at least one queue in the plurality of queues
2 comprises a first-in-first-out FIFO queue.

1 17. In a network interface apparatus, a method for managing transfer of data packets between
2 a host processor and a network, comprising:

3 managing memory in the network interface apparatus as a plurality of queues having
4 respective priorities, including placing a packet received from the host processor into one of the
5 plurality of queues according to a quality of service parameter associated with the packet, and
6 servicing packets in the plurality of queues according to the respective priorities; and
7 dynamically allocating space in said memory to the queues in the plurality of queues.

1 18. The method of claim 17, wherein the plurality of queues includes a higher priority queue,
2 and a lower priority queue, and including enabling a timeout timer coupled with the lower
3 priority queue if a packet is stored in the lower priority queue, the timeout timer expiring after a
4 timeout interval, and preempting the higher priority queue in favor of the lower priority queue if
5 the timeout timer expires.

1 19. The method of claim 17, wherein the plurality of queues includes a higher priority queue,
2 an intermediate priority queue, and a lower priority queue, and including

3 enabling a first timeout timer coupled with the intermediate priority queue if a packet is
4 stored in the intermediate priority queue, the first timeout timer expiring after a first timeout
5 interval, and preempting the higher priority queue in favor of the intermediate priority queue if
6 the first timeout timer expires; and

7 enabling a second timeout timer coupled with the lower priority queue if a packet is
8 stored in the lower priority queue, the second timeout timer expiring after a second timeout
9 interval, and preempting the higher priority queue and the intermediate priority queue in favor of
10 the lower priority queue if the second timeout timer expires.

1 20. The method of claim 17, wherein the plurality of queues includes a higher priority queue,
2 an intermediate priority queue, and a lower priority queue, and including

3 enabling a first timeout timer coupled with the intermediate priority queue if a packet is
4 stored in the intermediate priority queue, the first timeout timer expiring after a first timeout

interval, and preempting the higher priority queue in favor of the intermediate priority queue if the first timeout timer expires; and

enabling a second timeout timer coupled with the lower priority queue if a packet is stored in the lower priority queue, the second timeout timer expiring after a second timeout interval, and preempting the higher priority queue and the intermediate priority queue in favor of the lower priority queue if the second timeout timer expires; and

servicing the intermediate priority queue in favor of the lower priority queue if both the first and second timeout timers expire.

21. The method of claim 17, further comprising executing a security process on packets in one of the plurality of queues.

22. The method of claim 17, including formatting packets in the network interface device according to a protocol compliant with an Ethernet protocol standard.

23. The method of claim 17, including formatting packets in the network interface device packets according to a protocol compliant with an Infiniband protocol standard.

24. The method of claim 17, wherein the packets include frame start headers, and said quality of service parameters comprises codes in the frame start headers.

25. The method of claim 17, wherein said dynamically allocating space in said memory includes maintaining a list of free buffers and a list of used buffers for each of the plurality of queues.

26. The method of claim 17, wherein said memory comprises a single storage array.

27. The method of claim 17, wherein said dynamically allocating space in said memory includes placing packets downloaded to the memory in a plurality of buffers in non-contiguous memory locations.

1 28. The method of claim 17, wherein said dynamically allocating space in said memory
2 includes maintaining a list of free buffers and a list of used buffers for each of the plurality of
3 queues, so that each virtual path has a free buffer list having a number of free buffers, and
4 includes releasing a used buffer to the free buffer list for a queue in the plurality of queues
5 having a smallest number of free buffers.

1 29. The method of claim 17, wherein said dynamically allocating space in said memory
2 includes maintaining a list of free buffers and a list of used buffers for each of the plurality of
3 queues, so that each virtual path has a free buffer list having a number of free buffers, and
4 includes releasing a used buffer to the free buffer list for a queue in the plurality of queues
5 having a largest amount of traffic.

1 30. The method of claim 17, wherein said dynamically allocating space in said memory
2 includes maintaining a list of buffer descriptors for corresponding buffers in said memory, the
3 buffer descriptors including a parameter specifying a size of the corresponding buffer, and a
4 location of the corresponding buffer.

1 31. The method of claim 17, wherein said dynamically allocating space in said memory
2 includes maintaining a list of buffer descriptors for corresponding buffers in said memory, the
3 buffer descriptors including a programmable parameter specifying a size of the corresponding
4 buffer.

1 32. The method of claim 17, including managing at least one queue in the plurality of queues
2 as a first-in-first-out FIFO queue.

1 33. An integrated circuit for use in a network interface between a host processor and a
2 network, comprising:
3 a first port that receives data from the host processor;
4 a second port that transmits data to the network;
5 a memory that stores data packets received by the first port, the memory being coupled to
6 the first port and to the second port; and

7 a control circuit that manages the memory as a plurality of queues having respective
8 priorities, including logic to place a packet received on the first port into one of the plurality of
9 queues according to a quality of service parameter associated with the packet, and logic to
10 transmit packets in the plurality of queues out the second port according to the respective
11 priorities; and

12 logic to dynamically allocate space in said memory to the queues in the plurality of
13 queues.

1 34. The integrated circuit of claim 33, wherein the plurality of queues includes a higher
2 priority queue, and a lower priority queue, and including a timeout timer coupled with the lower
3 priority queue which is enabled if a packet is stored in the lower priority queue and expires after
4 a timeout interval, and including logic to preempt the higher priority queue in favor of the lower
5 priority queue if the timeout timer expires.

1 35. The integrated circuit of claim 33, wherein the plurality of queues includes a higher
2 priority queue, an intermediate priority queue, and a lower priority queue, and including
3 a first timeout timer coupled with the intermediate priority queue which is enabled if a
4 packet is stored in the intermediate priority queue and expires after a first timeout interval, and
5 including logic to preempt the higher priority queue in favor of the intermediate priority queue if
6 the first timeout timer expires; and

7 a second timeout timer coupled with the lower priority queue which is enabled if a packet
8 is stored in the lower priority queue and expires after a second timeout interval, and including
9 logic to preempt the higher priority queue and the intermediate priority queue in favor of the
10 lower priority queue if the second timeout timer expires.

1 36. The integrated circuit of claim 33, wherein the plurality of queues includes a higher
2 priority queue, an intermediate priority queue, and a lower priority queue, and including
3 an first timeout timer coupled with the intermediate priority queue which is enabled if a
4 packet is stored in the intermediate priority queue and expires after a first timeout interval, and
5 including logic to preempt the higher priority queue in favor of the intermediate priority queue if
6 the first timeout timer expires;

7 a second timeout timer coupled with the lower priority queue which is enabled if a packet
8 is stored in the lower priority queue and expires after a second timeout interval, and including
9 logic to preempt the higher priority queue and the intermediate priority queue in favor of the
10 lower priority queue if the second timeout timer expires; and

11 logic to service the intermediate priority queue in favor of the lower priority queue if both
12 the first and second timeout timers expire.

1 37. The integrated circuit of claim 33, further comprising logic in the network interface to
2 execute a security process on packets in one of the plurality of queues.

1 38. The integrated circuit of claim 33, wherein the second port further comprises media
2 access control circuitry for transmitting packets according to a protocol compliant with an
3 Ethernet protocol standard.

1 39. The integrated circuit of claim 33, wherein the second port further comprises media
2 access control circuitry for transmitting packets according to a protocol compliant with an
3 Infiniband protocol standard.

1 40. The integrated circuit of claim 33, wherein the packets include frame start headers, and
2 said quality of service parameters comprises codes in the frame start headers.

1 41. The integrated circuit of claim 33, wherein said logic to dynamically allocate space in
2 said memory maintains a list of free buffers and a list of used buffers for each of the plurality of
3 queues.

1 42. The integrated circuit of claim 33, wherein said memory comprises a single storage array.

1 43. The integrated circuit of claim 33, wherein said logic to dynamically allocate space in
2 said memory, places packets downloaded to the memory in a plurality of buffers in non-
3 contiguous memory locations.

1 44. The integrated circuit of claim 33, wherein said logic to dynamically allocate space in
2 said memory maintains a list of free buffers and a list of used buffers for each of the plurality of
3 queues, so that each virtual path has a free buffer list having a number of free buffers, and
4 includes logic which releases a used buffer to the free buffer list for a queue in the plurality of
5 queues having a smallest number of free buffers.

1 45. The integrated circuit of claim 33, wherein said logic to dynamically allocate space in
2 said memory maintains a list of free buffers and a list of used buffers for each of the plurality of
3 queues, so that each virtual path has a free buffer list having a number of free buffers, and
4 includes logic which releases a used buffer to the free buffer list for a queue in the plurality of
5 queues having a largest amount of traffic.

1 46. The integrated circuit of claim 33, wherein logic to dynamically allocate space in said
2 memory maintains a list of buffer descriptors for corresponding buffers in said memory, the
3 buffer descriptors including a parameter specifying a size of the corresponding buffer, and a
4 location of the corresponding buffer.

1 47. The integrated circuit of claim 33, wherein logic to dynamically allocate space in said
2 memory maintains a list of buffer descriptors for corresponding buffers in said memory, the
3 buffer descriptors including a programmable parameter specifying a size of the corresponding
4 buffer.

1 48. The integrated circuit of claim 33, wherein at least one queue in the plurality of queues
2 comprises a first-in-first-out FIFO queue.